arXiv:1702.01824v1 [stat.ML] 6 Feb 2017

Learning similarity preserving representations with neural similarity encoders

Franziska Horn & Klaus-Robert Müller
Machine Learning Group
Technische Universität Berlin
Berlin, Germany
franziska.horn@campus.tu-berlin.de
klaus-robert.mueller@tu-berlin.de

Abstract

Many dimensionality reduction or manifold learning algorithms optimize for retaining the pairwise similarities, distances, or local neighborhoods of data points. Spectral methods like Kernel PCA (kPCA) or isomap achieve this by computing the singular value decomposition (SVD) of some similarity matrix to obtain a low dimensional representation of the original data. However, this is computationally expensive if a lot of training examples are available and, additionally, representations for new (out-of-sample) data points can only be created when the similarities to the original training examples can be computed. We introduce similarity encoders (SimEc), which learn similarity preserving representations by using a feed-forward neural network to map data into an embedding space where the original similarities can be approximated linearly. The model optimizes the same objective as kPCA but in the process it learns a linear or non-linear embedding function (in the form of the tuned neural network), with which the representations of novel data points can be computed - even if the original pairwise similarities of the training set were generated by an unknown process such as human ratings. By creating embeddings for both image and text datasets, we demonstrate that SimEc can, on the one hand, reach the same solution as spectral methods, and, on the other hand, obtain meaningful embeddings from similarities based on human labels.

1 Introduction

Dimensionality reduction (DR) and manifold learning algorithms aim to represent data using fewer dimensions while preserving as much information about the original input data as possible. These condensed feature representations require less memory, they can often improve the performance in classification or regression tasks, for example, as they contain less noise, and by embedding the data in only two or three dimensions, a visualization of the dataset can be created for exploratory data analysis.

The variety of existing DR methods can be characterized based on several properties:

- The **objective** of the algorithm, i.e. what kind of information about the data should be preserved. This can either be the **variance** of the data, i.e. the lower dimensional representation should minimize the reconstruction error of the original data, or some **pairwise relations** between data points such as their similarities, distances, or nearest neighbors.
- Whether the transformation of the original data into the lower dimensional space is **linear or non-linear**.

- Whether the algorithm learns a function that can be used to **project new (out-of-sample) data points into the lower dimensional space** directly, or if only the representations for the given data points are learned.
- How the **optimization** is performed, i.e. whether the algorithm uses **gradient decent** to arrive at a (local) minima of the objective function or if the solution is convex, for example, obtained by performing a **singular value decomposition (SVD)** of the data's covariance or similarity matrix.

An overview of some common DR methods and how they compare in terms of these properties is given in Table 1.

Table 1: Dimensionality reduction methods.

	objective	complexity	out-of-sample solution	optimization
PCA [4, 7]	variance	linear	✓	SVD
Autoencoders [3, 12]	variance	non-linear	✓	gradient descent
Kernel PCA [10]	pairwise	non-linear	(✓) [1]	SVD
isomap [11]	pairwise	non-linear	(✓) [1]	SVD
LLE [9]	pairwise	non-linear	(✓) [1]	SVD
t-SNE [14, 15]	pairwise	non-linear	(✓) [13]	gradient descent
Similarity Encoder	pairwise	non-linear	✓	gradient descent

With the objective to retain the variance of the data, PCA [4, 7] and autoencoders [3, 12] are two great general purpose DR methods. If a linear solution is sufficient, PCA works well for a moderate number of original input dimensions, provided there are enough samples available to reliably estimate the covariance matrix.[1] If the patterns in the high dimensional input data are more complex, autoencoders can learn a non-linear transformation of the input data to the lower dimensional space. Since its objective function can be optimized using stochastic mini-batch gradient descent, the algorithm can easily learn from large datasets. Both PCA and autoencoders have the nice property that they learn (linear or non-linear) functions, with which out-of-sample (OOS) data points can be projected into the lower dimensional space directly.

Sometimes it is important to retain specific pairwise relations between the data points in the lower dimensional space. For example, t-SNE [15] is a very popular algorithm for creating two dimensional visualizations of datasets in which the local neighborhoods of data points are preserved. Great effort was put into extending this algorithm to work with large datasets [14] or to provide an explicit mapping function that can be applied to new data points [2, 13]. Other methods such as Kernel PCA (kPCA) [10], isomap [11], and locally linear embedding (LLE) [9] differ in their objective with respect to the type of similarity that should be preserved in the embedding, but they all obtain the lower dimensional coordinates of the input data points by performing the SVD of the corresponding similarity matrix. For large datasets and a correspondingly high dimensional similarity matrix, a SVD is computationally very expensive. Additionally, these methods can only embed OOS data points if their similarities to the original training samples can be computed [1]. This makes it impossible to create low dimensional representations of data points based on some arbitrary similarity matrices such as those obtained from human similarity judgments. Even if the given similarity matrix (approximately) abides by certain constraints, such as positive semi-definiteness, necessary to create a representation of the initially provided dataset [5], when new samples become available, these cannot be embedded into the lower dimensional space as we cannot compute their similarities to the existing examples (without consulting the human who provided the initial similarity judgments).

In this paper we introduce *similarity encoders* (SimEc), which learn similarity preserving representations by using a feed-forward neural network to map data into an embedding space where the original similarities can be approximated linearly. A SimEc can learn a linear or non-linear mapping function to project new data points into a lower dimensional embedding space. There are no constraints on the type of similarities that should be retained, i.e. mapping functions can even be learned for pairwise relations obtained from unknown similarity functions such as human ratings. Furthermore, the

[1] In case the number of feature dimensions is greater than the number of input samples, linear kernel PCA can be used to reach the same solution more efficiently [6].

algorithm can take advantage of large datasets for training since the objective function is optimized iteratively using stochastic mini-batch gradient descent. We show on both image and text datasets that SimEcs can, on the one hand, recreate solutions found by spectral methods such as kPCA or isomap, and, on the other hand, obtain meaningful embeddings from similarities based on human labels.

2 Similarity Encoders

We propose a novel dimensionality reduction framework termed similarity encoder (SimEc), which can be used to learn a linear or non-linear mapping function for computing low dimensional representations of data points such that the original pairwise similarities between the data points in the input space are preserved in the embedding space. For this, we borrow the "bottleneck" neural network (NN) architecture idea from autoencoders [3, 12]. Autoencoders aim to transform the high dimensional data points into low dimensional embeddings such that most of the data's variance is retained. Their network architecture has two parts: The first part of the network maps the data points from the original feature space to the low dimensional embedding (at the bottleneck). The second part of the NN mirrors the first part and projects the embedding back to a high dimensional output. This output is then compared to the original input to compute the reconstruction error of the training samples, which is used in the backpropagation procedure to tune the network's parameters. After the training is complete, i.e. the low dimensional embeddings encode enough information about the original input samples to allow for their reconstruction, the second part of the network is discarded and only the first part is used to project data points into the low dimensional embedding space. Similarity encoders have a similar two fold architecture, where in the first part of the network, the data is mapped to a low dimensional embedding, and then in the second part (which is again only used during training), the embedding is transformed such that the error of the representation can be computed. However, since here the objective is to retain the (non-linear) pairwise similarities instead of the data's variance, the second part of the NN does not mirror the first like it does in the autoencoder architecture.

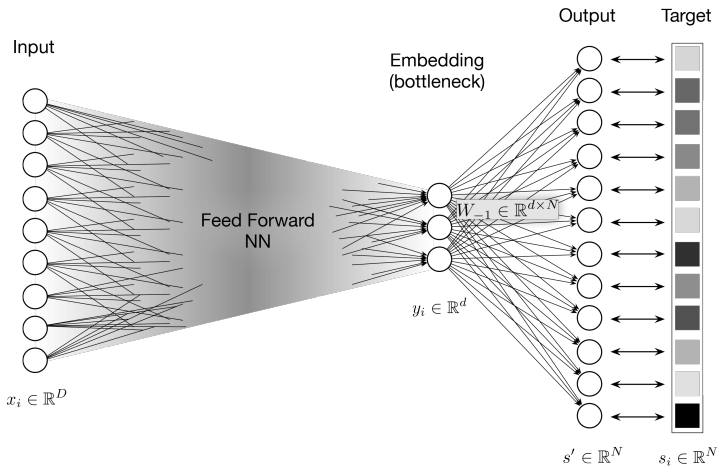

Figure 1: Similarity encoder (SimEc) architecture.

The similarity encoder architecture (Figure 1) uses as the first part of the network a flexible non-linear feed-forward neural network to map the high dimensional input data points $x_i \in \mathbb{R}^D$ to a low dimensional embedding $y_i \in \mathbb{R}^d$ (at the bottleneck). As we make no assumptions on the range of values the embedding can take, the last layer of the first part of the NN (i.e. the one resulting in the embedding) is always linear. For example, with two additional non-linear hidden layers, the embedding would be computed as

$$y_i = \sigma_1(\sigma_0(x_i W_0)W_1)W_2,$$

where σ_0 and σ_1 denote your choice of non-linear activation functions (e.g. tanh, sigmoid, or relu), but there is no non-linearity applied after multiplying with W_2. The second part of the network then consists of a single additional layer with the weight matrix $W_{-1} \in \mathbb{R}^{d \times N}$ to project the embedding

to the output, the approximated similarities $s' \in \mathbb{R}^N$:

$$s' = \sigma_{-1}(y_i W_{-1}).$$

These approximated similarities are then compared to the target similarities (for one data point this is the corresponding row $s_i \in \mathbb{R}^N$ of the similarity matrix $S \in \mathbb{R}^{N \times N}$ of the N training samples) and the computed error is used to tune the network's parameters with backpropagation. Applying the non-linearity σ_{-1} when computing the output is optional and it depends on the type of target similarities that the model is supposed to preserve whether this is advantageous or not.[2] As the cost function to minimize we choose the mean squared error between the output (approximated similarities) and the original (target) similarities. A regularization term is added to encourage the weights of the last layer (W_{-1}) to be orthogonal.[3] The model's objective function optimized during training is therefore:

$$\min \frac{1}{N} \sum_{i=1}^{N} \|s_i - s'\|_2^2 + \lambda \frac{1}{d^2 - d} \left\| W_{-1} W_{-1}^\top - \text{diag}(W_{-1} W_{-1}^\top) \right\|_1$$

where $\| \cdot \|_p$ denotes the respective p-norms for vectors and matrices and λ is a hyperparameter to control the strength of the regularization.

After the training is completed, only the first part of the neural network, which maps the input to the embedding, is used to create the representations of new data points. Depending on the complexity of the feed-forward NN, the mapping function learned by similarity encoders can be linear or non-linear, and because of the iterative optimization using stochastic mini-batch gradient descent, large amounts of data can be utilized to learn optimal representations.[4]

2.1 Relation to Kernel PCA

Kernel PCA (kPCA) is a popular non-linear dimensionality reduction algorithm, which performs the SVD of a kernel matrix to obtain low dimensional representations of the data points [10]. However, if the kernel matrix is very large this becomes computationally very expensive. Additionally, there are constraints on possible kernel functions (should be positive semi-definite) and new data points can only be embedded in the lower dimensional space if their kernel map (i.e. the similarities to the original training points) can be computed. As we show below, SimEc optimizes the same objective as kPCA but addresses these shortcomings.

The general idea is that both kPCA and SimEc embed the N data points in a feature space where the given target similarities can be approximated linearly (i.e. with the scalar product of the embedding vectors). When the error between the approximated (S') and the target similarities (S) is computed as the mean squared error, kPCA finds the optimal approximation by performing the SVD of the (centered) target similarity matrix, i.e.

$$S' = YY^\top,$$

where $Y \in \mathbb{R}^{N \times d}$ is the low dimensional embedding of the data based on the eigenvectors belonging to the d largest eigenvalues of S.

In addition to the embedding itself, it is often desired to have a parametrized mapping function, which can be used to project new (out-of-sample) data points into the embedding space. If the target similarity matrix is the linear kernel, i.e. $S = XX^\top$ where $X \in \mathbb{R}^{N \times D}$ is the given input data, this can easily be accomplished with traditional PCA. Here, the covariance matrix of the centered input data, i.e. $C = X^\top X$ is decomposed to obtain a matrix with parameters, $\tilde{W} \in \mathbb{R}^{D \times d}$, based on

[2] While a linear second part of the NN is useful to approximate the kPCA solution as illustrated in the next section, it can be beneficial to add a non-linearity on this final layer to approximate, for example, similarities based on class membership.

[3] To get embeddings similar to those obtained by kPCA, orthogonal weights in the last layer of the NN help as they correspond to the orthogonal eigenvectors of the kernel matrix found by kPCA.

[4] To speed up the training procedure and limit memory requirements for large datasets, the columns of the similarity matrix can also be subsampled (yielding $S \in \mathbb{R}^{N \times n}$), i.e. the number of target similarities (and the dimensionality of the output layer) is $n < N$, however all N training examples can still be used as input to train the network.

the eigenvectors belonging to the d largest eigenvalues of the covariance matrix. Then the optimal embedding (i.e. the same solution obtained by linear kPCA) can be computed as

$$Y = X\tilde{W}.$$

This serves as a mapping function, with which new data points can be easily projected into the lower dimensional embedding space.

When using a similarity encoder to embed data in a low dimensional space where the linear similarities are preserved, the SimEc's architecture would consist of a neural network with a single linear layer, i.e. the parameter matrix W_0, to project the input data X to the embedding $Y = XW_0$, and another matrix $W_{-1} \in \mathbb{R}^{d \times N}$ used to approximate the similarities as

$$S' = YW_{-1}.$$

From these formulas one can immediately see the link between linear similarity encoders and PCA / linear kPCA: once the parameters of the neural network are tuned correctly, W_0 would correspond to the mapping matrix $\tilde{W}$ found by PCA and W_{-1} could be interpreted as $Y^\top$, i.e. Y would be the same eigenvector based embedding as found with linear kPCA.

Finding the corresponding function to map new data points into the embedding space is trivial for linear kPCA, but this is not the case for other kernel functions. While it is still possible to find the optimal embedding with kPCA for non-linear kernel functions, the mapping function remains unknown and new data points can only be projected into the embedding space if we can compute their kernel map, i.e. the similarities to the original training examples [1]. Some attempts were made to manually define an explicit mapping function to represent data points in the kernel feature space, however this only works for specific kernels and there exists no general solution [8]. As neural networks are universal function approximators, with the right architecture similarity encoders could instead learn arbitrary mapping functions for unknown similarities to arrive at data driven kernel learning solutions.

2.2 Model overview

The properties of similarity encoders are summarized in the following. The objective of this dimensionality reduction approach is to retain pairwise similarities between data points in the embedding space. This is achieved by tuning the parameters of a neural network to obtain a linear or non-linear mapping (depending on the network's architecture) from the high dimensional input to the low dimensional embedding space where the target similarities are approximated linearly. Since the cost function is optimized using stochastic mini-batch gradient descent, we can take advantage of large datasets for training. The embedding for new test points can be easily computed with the explicit mapping function in the form of the tuned neural network. And since there is no need to compute the similarity of new test examples to the original training data for out-of-sample solutions (like with kPCA), the target similarities can be generated by an unknown process such as human similarity judgments.

3 Experiments

In the following experiments we demonstrate that similarity encoders can, on the one hand, reach the same solution as kPCA, and, on the other hand, generate meaningful embeddings from human labels. To illustrate that this is independent of the type of data, we present results obtained both on the well known MNIST handwritten digits dataset as well as the 20 newsgroups text corpus. Further details as well as the code to replicate these experiments and more is available online.[5]

We compare the embedding found with linear kPCA to that created with a linear similarity encoder (consisting of one linear layer mapping the input to the embedding and a second linear layer to project the embedding to the output, i.e. computing the approximated similarities). Additionally, we show that a non-linear SimEc can approximate the solution found with isomap (i.e. the SVD of the geodesic distance matrix). To quantify how well the SimEc solution approximates the low dimensional representations created with the respective spectral method independent of rotation and scaling, we compute the Pearson correlation coefficient between the normalized pairwise euclidean

[5] `https://github.com/cod3licious/simec/examples_simec.ipynb`

distances of all points in the embedding spaces. We found that for optimal results the kernel matrix used as the target similarity matrix for the SimEc should first be centered (as it is being done for kPCA as well [6]).

In a second step, we show that SimEcs can learn the mapping to a low dimensional embedding for arbitrary similarity functions and reliably create representations for new test samples without the need to compute their similarities to the original training examples, thereby going beyond the capabilities of kPCA. For both datasets we illustrate this by using the class labels assigned to the samples by human annotators to create the target similarity matrix for the training fold of the data, i.e. S is 1 for data points belonging to the same class and 0 everywhere else. In this setup, the similarity matrix is not centered but instead, since the targets are binary, we use a non-lineaer activation function in the final layer to ensure the predicted similarities are between 0 and 1 as well:[6]

$$s' = \sigma_{-1}(y_i W_{-1}) \quad \text{with} \quad \sigma_{-1}(z) = \frac{1}{1 + e^{-10(z-0.5)}}.$$

We compare the solutions found by SimEc architectures with a varying number of additional non-linear hidden layers in the first part of the network (while keeping the embedding layer linear as before) to show how a more complex network improves the ability to map the data into an embedding space in which the class-based similarities are retained.

MNIST The MNIST dataset contains 28×28 pixel images depicting handwritten digits. For our experiments we randomly subsampled 10k images from all classes, of which 80% are assigned to the training fold and the remaining 20% to the test fold (in the following plots, data points belonging to the training set are displayed transparently while the test points are opaque). As shown in Figure 2, the embeddings of the MNIST dataset created with linear kPCA and a linear similarity encoder, which uses as target similarities the linear kernel matrix, are almost identical (up to a rotation and scaling). The same holds true for the isomap embedding, which is well approximated by a non-linear SimEc with two hidden layers using the geodesic distances between the data points as targets (Figure 5 in the Appendix). When optimizing SimEcs to retain the class-based similarities (Figure 3), additional

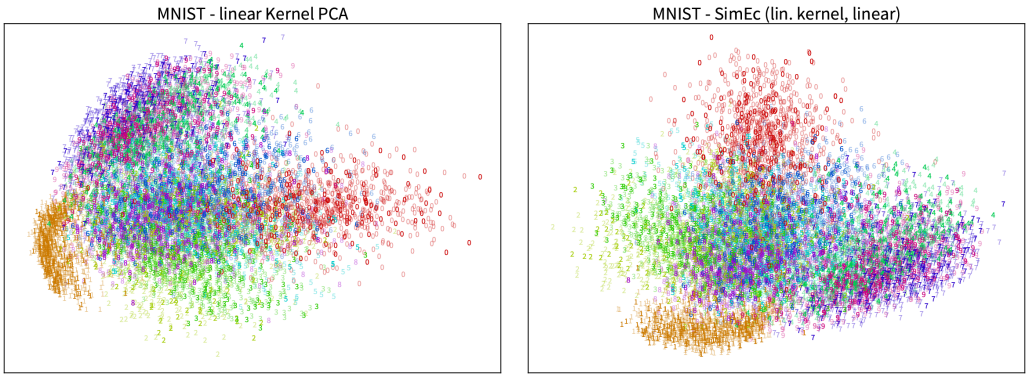

Figure 2: MNIST digits visualized in two dimensions by linear kPCA and a linear SimEc ($r = 0.96$).

non-linear hidden layers in the feed-forward NN can improve the embedding by further separating data points belonging to different classes in tight clusters. As it can be seen, the test points (opaque) are nicely mapped into the same locations as the corresponding training points (transparent), i.e. the model learns to associate the input pixels with the class clusters only based on the imposed similarities between the training data points.

20 newsgroups The 20 newsgroups dataset consists of around 18k newsgroup posts assigned to 20 different topics. We take a subset of seven categories and use the original train/test split ($\sim$4.1k and $\sim$2.7k samples respectively) and remove metadata such as headers to avoid overfitting.[7] All text documents are transformed into 46k dimensional tf-idf feature vectors, which are used as input to

[6] This scaled and shifted sigmoid function maps values between 0 and 1 almost linearly while thresholding values outside this interval.

[7] `http://scikit-learn.org/stable/datasets/twenty_newsgroups.html`

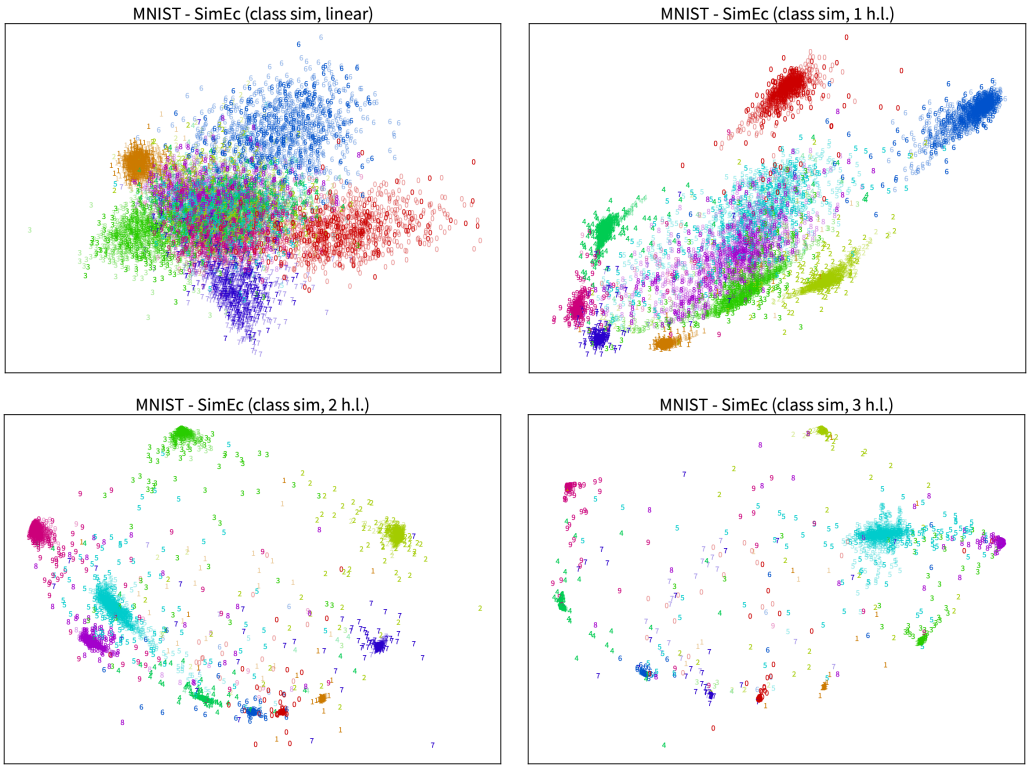

Figure 3: MNIST digits visualized in two dimensions by SimEcs with an increasing number of non-linear hidden layers and the objective to retain similarities based on class membership.

the SimEc and to compute the linear kernel matrix of the training fold. The embedding created with linear kPCA is again well approximated by the solution found with a corresponding linear SimEc (Figure 6 in the Appendix). Additionally, this serves as an example where traditional PCA is not an option to obtain the corresponding mapping matrix for the linear kPCA solution, as due to the high dimensionality of the input data and comparatively low number of samples, the empirical covariance matrix would be poorly estimated and too large to decompose into eigenvalues and -vectors. With the objective to retain the class-based similarities, a SimEc with a non-linear hidden layer clusters documents by their topics (Figure 4).

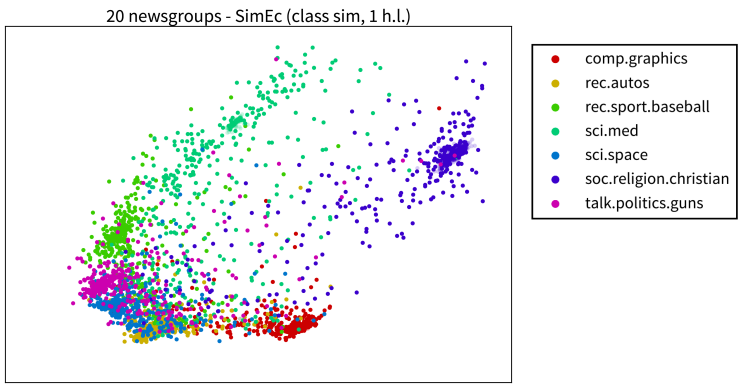

Figure 4: 20 newsgroups texts visualized in two dimensions by a non-linear SimEc with one hidden layer and the objective to preserve the similarities based on class membership in the embedding.

4 Discussion & Outlook

Representing intrinsically complex data is an ubiquitous challenge in data analysis. While kernel methods and manifold learning have made very successful contributions, their ability to scale is somewhat limited. Neural autoencoders offer scalable nonlinear embeddings, but their objective is to minimize the reconstruction error of the input data, which does not necessarily preserve important pairwise relations between data points. In this paper we have proposed SimEcs as a neural network framework which bridges this gap by optimizing the same objective as spectral methods, such as kPCA, for creating similarity preserving embeddings while retaining the favorable properties of autoencoders.

Similarity encoders are a novel method to learn similarity preserving embeddings and can be especially useful when it is computationally infeasible to perform the SVD of a kernel matrix, when the target similarities are obtained through an unknown process such as human similarity judgments, or when an explicit mapping function is required. To accomplish this, a feed-forward neural network is constructed to map the data into an embedding space where the original similarities can be approximated linearly. We have demonstrated the usefulness of SimEcs by visualizing data from different domains, going beyond the capabilities of traditional methods.

Future work will aim to further the theoretical understanding of SimEcs and explore other application scenarios where using this novel neural network architecture can be beneficial. As it is often the case with neural network models, determining the optimal architecture as well as other hyperparameter choices best suited for the task at hand can be difficult. While so far we mainly studied SimEcs based on fairly simple feed-forward networks, it appears promising to consider also deeper neural networks and possibly even more elaborate architectures, such as convolutional networks, for the initial mapping step to the embedding space, as in this manner hierarchical structures in complex data could be reflected. Note furthermore that prior knowledge as well as more general error functions could be employed to tailor the embedding to the desired application target(s).

Acknowledgments

We would like to thank Antje Relitz, Christoph Hartmann, and Ivana Balažević for their helpful comments on earlier versions of this manuscript.
Franziska Horn acknowledges funding from the Elsa-Neumann scholarship from the TU Berlin.

References

[1] Yoshua Bengio, Jean-François Paiement, Pascal Vincent, Olivier Delalleau, Nicolas Le Roux, and Marie Ouimet. Out-of-sample extensions for LLE, Isomap, MDS, Eigenmaps, and Spectral Clustering. *Advances in neural information processing systems*, 16:177–184, 2004.

[2] Kerstin Bunte, Michael Biehl, and Barbara Hammer. A general framework for dimensionality-reducing data visualization mapping. *Neural Computation*, 24(3):771–804, 2012.

[3] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006.

[4] Ian Jolliffe. *Principal component analysis*. Wiley Online Library, 2002.

[5] Julian Laub and Klaus-Robert Müller. Feature discovery in non-metric pairwise data. *The Journal of Machine Learning Research*, 5:801–818, 2004.

[6] Klaus-Robert Müller, Sebastian Mika, Gunnar Rätsch, Koji Tsuda, and Bernhard Schölkopf. An introduction to kernel-based learning algorithms. *Neural Networks, IEEE Transactions on*, 12(2):181–201, 2001.

[7] Karl Pearson. On lines and planes of closest fit to systems of point in space. *Philosophical Magazine*, 2 (11):559–572, 1901.

[8] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2007.

[9] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[10] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural computation*, 10(5):1299–1319, 1998.

[11] Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[12] Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
[13] Laurens van der Maaten. Learning a parametric embedding by preserving local structure. In *International Conference on Artificial Intelligence and Statistics*, pages 384–391, 2009.
[14] Laurens van der Maaten. Barnes-Hut-SNE. In *Proceedings of the International Conference on Learning Representations*, 2013.
[15] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(2579-2605):85, 2008.

Appendix

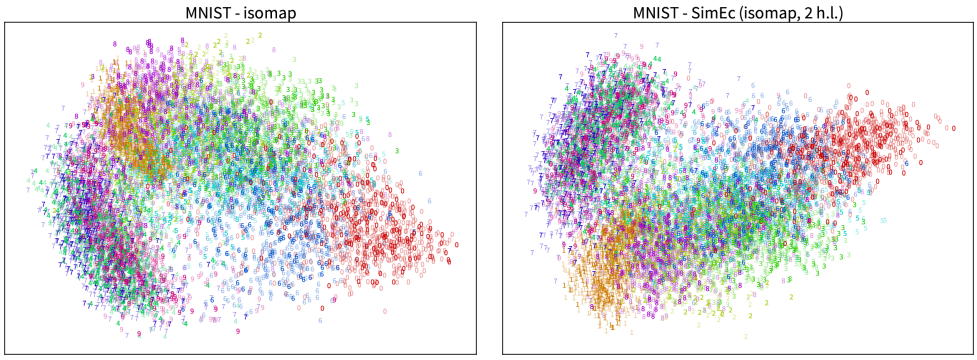

Figure 5: MNIST digits visualized in two dimensions by isomap and a non-linear SimEc ($r = 0.92$).

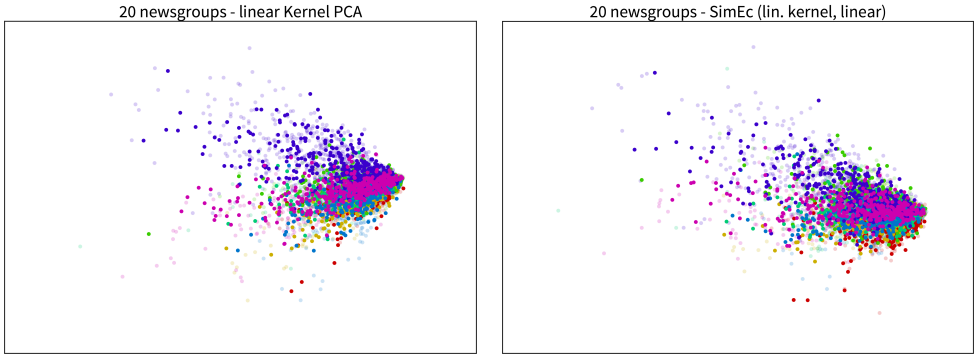

Figure 6: 20 newsgroups dataset embedded with linear kPCA and a corresponding linear SimEc ($r = 0.96$).